

Persistence of reduced neuromotor noise in long-term motor skill learning

Meghan E. Huber,¹ Nikita Kuznetsov,² and Dagmar Sternad^{2,3,4,5}

¹Department of Bioengineering, Northeastern University, Boston, Massachusetts; ²Department of Biology, Northeastern University, Boston, Massachusetts; ³Department of Electrical and Computer Engineering Northeastern University, Boston, Massachusetts; ⁴Department of Physics, Northeastern University, Boston, Massachusetts; and ⁵Center for the Interdisciplinary Research on Complex Systems, Northeastern University, Boston, Massachusetts

Submitted 29 March 2016; accepted in final form 26 September 2016

Huber ME, Kuznetsov N, Sternad D. Persistence of reduced neuromotor noise in long-term motor skill learning. *J Neurophysiol* 116: 2922–2935, 2016. First published September 28, 2016; doi:10.1152/jn.00263.2016.—It is well documented that variability in motor performance decreases with practice, yet the neural and computational mechanisms that underlie this decline, particularly during long-term practice, are little understood. Decreasing variability is frequently examined in terms of error corrections from one trial to the next. However, the ubiquitous noise from all levels of the sensorimotor system is also a significant contributor to overt variability. While neuromotor noise is typically assumed and modeled as immune to practice, the present study challenged this notion. We investigated the long-term practice of a novel motor skill to test whether neuromotor noise can be attenuated, specifically when aided by reward. Results showed that both reward and self-guided practice over 11 days improved behavior by decreasing noise rather than effective error corrections. When the challenge for obtaining reward increased, subjects reduced noise even further. Importantly, when task demands were relaxed again, this reduced level of noise persisted for 5 days. A stochastic learning model replicated both the attenuation and persistence of noise by scaling the noise amplitude as a function of reward. More insight into variability and intrinsic noise and its malleability has implications for training and rehabilitation interventions.

motor learning; neuromotor noise; retention; skill acquisition; variability

NEW & NOTEWORTHY

To date, the question of whether neuromotor noise is regulated with practice, especially during long-term practice, has not been addressed. Moreover, it is unclear whether noise in neuromotor output can be modulated with lasting effects. The results of this study demonstrate that complex skill learning proceeds by reducing neuromotor noise, suggesting that error correction is not the only mechanism for improving motor performance. Furthermore, our results reveal how noise reduction can be expedited and retained.

ACQUIRING A NEW MOTOR SKILL requires long hours of practice and patience, regardless of whether the goal is to play for the Boston Bruins or merely to make the high school hockey team. The same patience is required when recovering from brain injury such as stroke, where many months must be dedicated to relearning seemingly basic motor skills (Jørgensen et al. 1995). And yet, when motor learning is studied in the laboratory,

practice sessions rarely exceed a single hour, for practical reasons. Despite several experimental documentations of long-term practice and retention (Cohen and Sternad 2009, 2012; Nourrit-Lucas et al. 2013; Park et al. 2013; Park and Sternad 2015), little is known about what physiological or computational mechanisms bring about the fine-tuning and persistence of the acquired skill. The present study examined 1) what processes may underlie the long-term improvements and retention and 2) how such honing of skill can be enhanced.

Skill learning is marked by decrements in error and variability in behavior, starting with relatively rapid changes and followed by subtle tuning that continues over weeks, if not years, of practice (Crossman 1959). In short-term motor adaptation studies, reduction in error and variability with practice has been attributed to the amount of error corrected from one trial to the next, i.e., the correction gain (Herzfeld et al. 2014; Smith et al. 2006; Thoroughman and Shadmehr 2000). Yet once the error correction gain has been optimized and errors have become very small, how the observed variability continues to decrease over long-term practice is still an open question. Many previous studies on variability have examined simple tasks, such as reaching to a target, where improvements are only visible over the first few trials as the task is already well practiced (van Beers et al. 2013). Other studies have examined learning of more complex novel tasks, often by comparing novices and experts, where variability is frequently quantified in terms of standard deviations, assuming a Gaussian distribution, without parsing of its potential structure (Schmidt and Lee 2011). The present study aimed to examine overt variability over long practice and to identify structure due to error corrections and random noise processes.

Random fluctuations in behavioral data, seen both in continuous trajectories and in trial-to-trial variability, are often ascribed to neuromotor noise, i.e., the noise that arises from the many interconnected layers of deterministic and stochastic processes throughout the neuromotor system, such as that from variations in synaptic transmission and motor unit recruitment (Faisal et al. 2008). Even though prior studies have shown that the magnitude of variability and the implied underlying neuromotor noise can depend on movement speed (Fitts 1954; Woodworth 1899) and signal strength (Harris and Wolpert 1998), the noise in the output of the neuromotor system is commonly viewed and modeled as immune to practice (e.g., van Beers 2009; van Beers et al. 2013). The present study challenged this assumption. We hypothesized that in long-term skill learning the variance of neuromotor noise decreases to attenuate overt variability.

Address for reprint requests and other correspondence: M. E. Huber, Massachusetts Inst. of Technology, Dept. of Mechanical Engineering, Bldg. 3, Rm. 143, 77 Massachusetts Ave., Cambridge, MA 02139 (e-mail: mehuber@mit.edu).

If neuromotor noise can decrease, then augmented feedback known to facilitate learning should be able to enhance this effect. Reward or “satisfying” consequences, recognized as a mediator for learning for over a century (Thorndike 1927), have received renewed attention with supporting evidence in motor adaptation (Abe et al. 2011; Galea et al. 2015; Izawa and Shadmehr 2011; Nikooyan and Ahmed 2015; Shmuelof et al. 2012). Grounded in this rich history and recent evidence for reinforcement learning, we used reward to determine the extent to which neuromotor noise could be diminished.

We performed two experiments to test the overall hypothesis that the decrease in behavioral variability observed in novel skill learning proceeds by reducing neuromotor noise, especially in later stages of practice. In each experiment, we assessed learning of a virtual throwing task over 11 daily sessions. In *experiment 1*, we hypothesized that reward leads to faster learning and better performance, both in error and variability, compared with self-guided learning (*hypothesis 1*). Furthermore, we expected that with extended practice the fine-tuning of skill is achieved by decreasing the variance or amplitude of noise (*hypothesis 2*). Initial experimental results showed that while reward accelerated improvement, the reward group and the self-guided control group reached the same level of performance and noise after 11 practice days. Hence, in *experiment 2*, we hypothesized that decreasing the reward would lead subjects to decrease their variability through decreasing noise variance further (*hypothesis 3*). We additionally expected that the low level of neuromotor noise would persist for 5 days, even after the reward returned to its initial level, indicating long-term retention of the skill (*hypothesis 4*). The overall premise is that overt variability during learning arises from two main sources: error corrections and noise from the neuromotor system. Our analyses demonstrate that error corrections did not play a measurable role. Hence, we conclude that the overt variability is due to noise from the neuromotor system. A simple iterative model with a noise source that depended on reward was used to explain these experimental findings including retention.

MATERIALS AND METHODS

Participants

Eighteen healthy right-handed students (9 women and 9 men, mean age 25.1 ± 2.3 yr) from Northeastern University took part in the two experiments. None had any prior experience with the experimental task. Each participant performed 240 throws per day of a virtual throwing task for 11 days, resulting in a total of 2,640 throws. All participants gave informed written consent before the experiment and received monetary compensation upon completion of all 11 daily sessions. The experimental protocol was reviewed and approved by the Institutional Review Board of Northeastern University.

Experimental Task and Apparatus

The experimental task was modeled after the British pub game skittles. In skittles, the player throws a ball tethered to a center post to hit a target skittle at the far side of the post (Fig. 1A). This task is one variant of throwing, which has received much attention in motor control as it requires precise timing of the ball release to achieve hitting accuracy (Cohen and Sternad 2012; Hore et al. 1995; Smeets et al. 2002). In the present experimental variant, the participant manipulated a horizontal lever arm with a single-joint rotation about

the elbow to throw a ball to hit a target in a two-dimensional virtual environment (Fig. 1, B and C; see also Cohen and Sternad 2009).

A two-dimensional model in which the ball was attached to the origin by two orthogonal massless springs determined the path of the ball upon release. Because of restoring forces proportional to the distance of the ball from the origin and the velocity imparted at release, the ball was accelerated around the origin, traversing an elliptical trajectory. The equations for ball position in the x - and y -directions at time t were

$$x(t) = A_x \sin(\omega t + \varphi_x) e^{-\frac{t}{\tau}}$$

$$y(t) = A_y \sin(\omega t + \varphi_y) e^{-\frac{t}{\tau}}$$

The frequency ω denotes the natural frequency of the system, and relaxation time τ was used to introduce damping to approximate realistic behavior (for more details on the physical model, see Hasson et al. 2016). For this study, $\omega = 3.16$ rad/s and $\tau = 20$ s. Intuitively, the frequency ω determined the speed at which the ball traversed the post; the relaxation time τ determined the energy loss or how much the ball trajectory spiraled toward the equilibrium position of the post. The amplitudes A_x and A_y and the phases φ_x and φ_y of the sinusoidal motions of the two springs were calculated from the angular position and velocity of the ball at release based on the recorded movement of the manipulandum.

A top-down view of the skittles workspace was displayed on a rear projection screen ~ 150 cm in front of the participant. Figure 1C shows the virtual scene, which consisted of the post (radius = 25 cm) centered at the origin, the target (radius = 5 cm) located 105.5 cm above and 5 cm to the right of the origin, the ball (radius = 5 cm), and the lever arm (length = 40 cm), with the axle located 150 cm below the origin. The positions and size of each object were defined in the virtual coordinate system and then scaled to project the graphics on the large screen. The gain between the real projected and the virtual workspace was 2.27, such that real radius of the ball and target on the projection screen was 2.20 cm and the post radius was 11 cm. Taking the distance of the subject's viewpoint into account, the visual angle of the projected objects was thus calculated for the post (8.39°) and the target and ball (1.67°). Assuming normal visual acuity, subjects should have had no problem discerning the distance from the ball trajectory to the target center.

Participants moved a real lever arm to control the angle of the virtual arm. The participant's forearm was secured to the horizontal manipulandum that rotated around an axle centered at the elbow joint. A potentiometer (Vishay Spectrol, Ontario, CA) at the axle continuously sampled the angular position of the manipulandum at ~ 700 Hz with a DT300 data acquisition board (Data Translation, Marlborough, MA) into a custom C++ program that ran the simulation.

Participants were instructed to release the ball such that it traveled through the center of the target without hitting the post. At the start of each throw, the participant grasped a wooden ball affixed to the distal end of manipulandum and closed a contact switch with his or her index finger (Trossen Robotics, Westchester, IL). This attached the virtual ball to the end of the virtual arm. The participant then moved his or her arm and triggered the ball release by extending the index finger to open the switch. Upon release, all participants saw the ball traverse an elliptical path in the virtual scene for 1.5 s with a shape determined by the specific angle and velocity of the manipulandum at the moment of release. The elliptical shape of the ball trajectory also presented important information to the subjects about their ball release and hitting success. Furthermore, after the ball stopped, a short segment of the ball trajectory close to the target center provided additional visual feedback.

Error for a given throw was defined as the minimum distance between the ball path and the target center (Fig. 1D). In the experimental conditions in which subjects were rewarded for successful throws, the target turned from yellow to green when a throw resulted

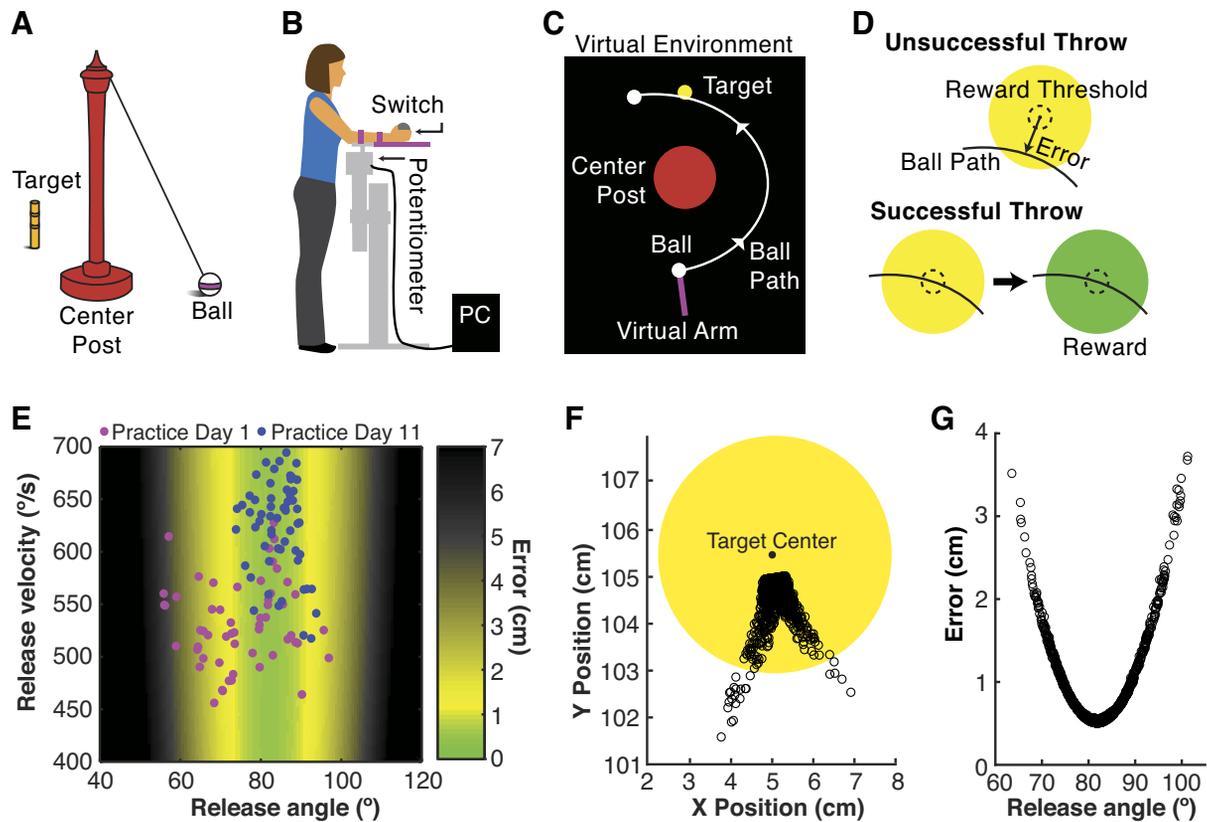


Fig. 1. Task description for the skill learning experiments. *A*: real skittles throwing task. The experimental task was modeled after the British pub game skittles. In skittles, the player throws a ball tethered to the top of a post around that post to hit a target skittle on the other side. *B* and *C*: experimental setup of virtual throwing task. Subjects manipulated a horizontal lever arm with a single-joint rotation about the elbow to throw a virtual ball in a 2-dimensional virtual environment as seen in *C*. As in the real game, they were instructed to throw the ball such that it traveled through the center of the target without hitting the post. The origin of the x - y workspace was defined at the center of the post. The target was located at $x = 5$ cm, $y = 105$ cm. *D*: error was defined as the minimum distance between the ball path and the target center. If the error was smaller than the reward threshold, the subject received a reward indicating a successful target hit; otherwise there was no added feedback. The reward was given in the form of a color change of the target from yellow to green. Subjects in the reward and changing-reward groups were instructed to achieve as many target hits as possible. Note that subjects in the self-guided group were instructed to hit the target center as accurately as possible, and the target always remained yellow, regardless of error. *E*: result space of modeled skittles task. The result space illustrates the functional relation between the 2 variables that determine the ball path, release angle and velocity, and the result variable, error: for each combination of release angle and velocity, error is depicted by color. The dots depict 60 throws (represented by their release angle-velocity pairs) of an example subject in the reward group during the first (magenta) and last (blue) days of practice. *F*: location of individual throws in x - y workspace at which the ball came closest to the target, visualized for 1 example subject over 3 days (720 throws). *G*: the relation between release angle and error is nonlinear, showing the redundancy: any given error can be achieved with >1 release angle.

in an error smaller than the defined reward threshold, indicating a successful hit to the subject. Note that the target size remained the same in all conditions, and the size of the reward threshold was not visible on the screen.

Since the two execution variables, angle and velocity at ball release, fully determined the ball trajectory, they also fully determined the result variable, error. The result space in Fig. 1*E* illustrates this functional relationship: for each point in execution space, defined by release angle and velocity, the error for that throw is depicted by color: green defines successful (i.e., rewarded) solutions; yellow to black shades show increasing error. Successful hits of the target (minimum error) defined the solution manifold.

The target location chosen for this study simplified the task such that release velocity did not affect the result; only the arm angle at ball release defined performance success. A previous study has shown that humans have a remarkably good proprioceptive sense to reproduce a given joint angle pose (Cordo 1990). Note that different target locations lead to very different solution manifolds (Sternad et al. 2014). This target configuration and solution space made the task faster to learn, turning the focus on the refinement of motor skill rather than the initial exploratory stage. Furthermore, as release velocity did not contribute to hitting success, covariation between release angle

and velocity could not be exploited to reduce the variability in performance. To reduce variability in performance subjects could only reduce variability in release angle.

Experiment 1: Self-Guided vs. Reward Learning

Experiment 1 compared two groups that practiced with different conditions: the reward group ($n = 6$) received reward in the form of a color change when the error was below a threshold. The reward threshold was set to 1.1 cm in the virtual coordinates (0.48 cm on the projection screen, which corresponded to a visual angle of 0.37°). The self-guided group ($n = 6$) practiced without receiving any reward for successful trials.

Experiment 2: Changing Reward Threshold During Learning

Experiment 2 examined how the threshold for the reward signal influenced the performance. Hence, a changing-reward group of participants ($n = 6$) practiced the task with different reward thresholds throughout practice. In the baseline block (days 1–3), the reward threshold was 1.1 cm as in *experiment 1*; in the manipulation block (days 4–6), the threshold was changed to 0.65 cm (corresponding to

0.29 cm in the projected workspace with a visual angle of 0.22° , making it more difficult to receive a reward; in the persistence block (*days 7–11*), the initial reward threshold of 1.1 cm was reinstated, making reward again easier to obtain. To assess the influence of the reward threshold, performance of the changing-reward group was compared to that of the group with constant reward from *experiment 1*.

Dependent Measures

Because of the specific target location, the ball trajectory could not go through the target center (zero error), nor could it go around, i.e., overshoot the target. Hence, the nearest point that determined the minimum distance error was below the target, as shown in exemplary data from 720 trials of one subject (Fig. 1*F*). For the same data, Fig. 1*G* illustrates that release angles map into error with an approximately parabolic function; note that two release angles could achieve the same error. Furthermore, even if the data were symmetrically distributed around the optimal release angle, the nonnegative definition of error resulted necessarily in a skewed distribution of error. Therefore, we used median error as one performance measure. Because the error was unsigned (i.e., a measure of absolute deviation from the target center), the median error on each practice day reflected the combined effects of bias and variability in performance (Howell 2014). For the two reward groups, task performance was also evaluated by the percentage of rewarded throws (errors smaller than the reward threshold) of each practice day.

To measure variability in motor execution, we turned to variability of release angle as metric. Unlike the error, there is no redundancy in the release angle, which allows for a more direct and accurate measure of variability in the neuromotor output. Because of the frequent non-Gaussian distribution (in 40% of the data sets), interquartile range (IQR) served as metric of dispersion.

As Fig. 1*E* shows, the solution manifold of this target location was parallel to velocity, meaning that the success of the throw was independent of velocity. Therefore, motor execution was fully quantified by release angles and error. Nevertheless, to determine whether reductions in release angle variability were possibly achieved by movements with slower velocity, the median release velocity on each day was also analyzed.

In addition to this distributional analysis, the temporal structure of the trial-by-trial sequence was examined by calculating the lag-1 autocorrelation coefficient (ACF-1) of the release angles on each day. This measure was used to detect the presence of correlations between successive throws. ACF-1 values between -1 and 0 indicate that successive observations were negatively correlated, also known as antipersistence; ACF-1 values between 0 and 1 indicate persistence, or positive correlation between successive observations. An ACF-1 value of 0 would indicate that the signal is uncorrelated white noise. The 95% confidence interval for the ACF-1 of a white noise process is $\pm 1.96/\sqrt{N}$, where N is the length of the signal. Thus observing an ACF-1 value between approximately ± 0.126 for the release angles on a given day ($n = 240$) would suggest that the signal is a white noise process.

Detrended fluctuation analysis (DFA) was also applied to the time series of release angles to identify temporal dependencies beyond lag-1 (Abe and Sternad 2013; Dingwell and Cusumano 2010; Peng et al. 1994). Importantly, DFA can identify persistence and antipersistence without exact sequential dependencies. The scaling index quantifies the long-range correlations of the time series. Here it was assessed for the timescales of 5–20 throws to complement the results of ACF-1. Scaling indexes above and below 0.5 indicate correlation and anticorrelation, respectively; a scaling index of 0.5 indicates lack of any detectable structure, i.e., noise.

Statistical Analyses

The dependent measures were analyzed for the 240 throws on each day of practice. The median and IQR of the dependent measures were used because the Shapiro-Wilk tests revealed that the distributions of each measure were not normal on all days (Shapiro and Wilk 1965). In each experiment, the dependent measures were analyzed with two-way repeated-measures ANOVAs (Group \times Day). Group was a between-participants factor, and Day was a within-participants factor. The Greenhouse-Geisser correction factor was applied to the within-subject effects when sphericity was violated (Kirk 1995). Relevant planned comparisons using paired *t*-tests were made to further probe the within-participant effects; independent-sample *t*-tests investigated group effects. While the ANOVAs were used to assess group differences during practice, independent-sample *t*-tests on *days 10* and *11* were used to directly assess group difference in skill level at the end of practice. The significance level was set at $\alpha = 0.05$ for all statistical tests.

RESULTS

Reward Accelerates Error Reduction, But After Extended Practice Self-Guided Practice Reaches Same Level of Performance

Experiment 1 tested *hypothesis 1* stating that subjects who were incentivized by reward would reduce both their error and variability faster and to a lower level. Subjects in the reward group ($n = 6$) received a visual signal of success after each throw when the ball trajectory passed through the target with an error < 1.1 cm (Fig. 1*D*). The self-guided group ($n = 6$) performed the task without any reward signal, and the target remained yellow in all cases.

As expected, subjects in both groups reduced their median error over 11 days of practice, as seen in a significant effect of day [$F_{(10,100)} = 9.57, P = 0.002$, partial $\eta^2 = 0.49$] (Fig. 2*A*). To determine whether subjects reached an asymptote in their overall performance, paired *t*-tests compared median error on each day to the last day of practice. From *day 4* onward, there was no significant change in median error compared with *day 11* (all $P > 0.10$), suggesting that subjects did reach an asymptote in performance by the end of practice. The reward group also increased their percentage of rewarded throws from an initial average of 69% on *day 1* to 84% by *day 11*. Subjects in the reward group performed with lower median error than the self-guided group, indicated by a significant group effect [$F_{(1,10)} = 15.62, P = 0.003$, partial $\eta^2 = 0.61$]. The interaction was also significant [$F_{(10,100)} = 4.34, P = 0.034$, partial $\eta^2 = 0.30$], because the self-guided group began with greater median error on the first 2 days ($P = 0.001$ and $P = 0.0046$) but approached the reward group on most of the following days (see Fig. 2*A*). This interaction suggests that reward accelerates error reduction as hypothesized. Independent *t*-tests detail this interaction as shown in Fig. 2*A*: most important to *hypothesis 1* is that *t*-tests between the self-guided and reward groups failed to detect a statistically significant difference on *day 10* [$t(10) = 1.12, P = 0.29, d = 0.65$] and *day 11* [$t(10) = 2.13, P = 0.06, d = 1.23$]. While these results may be underpowered because of the small sample size, the convergence on both *days 10* and *11* suggested that with sufficient practice the self-guided group reached the same level of error, a result that ran counter to *hypothesis 1*.

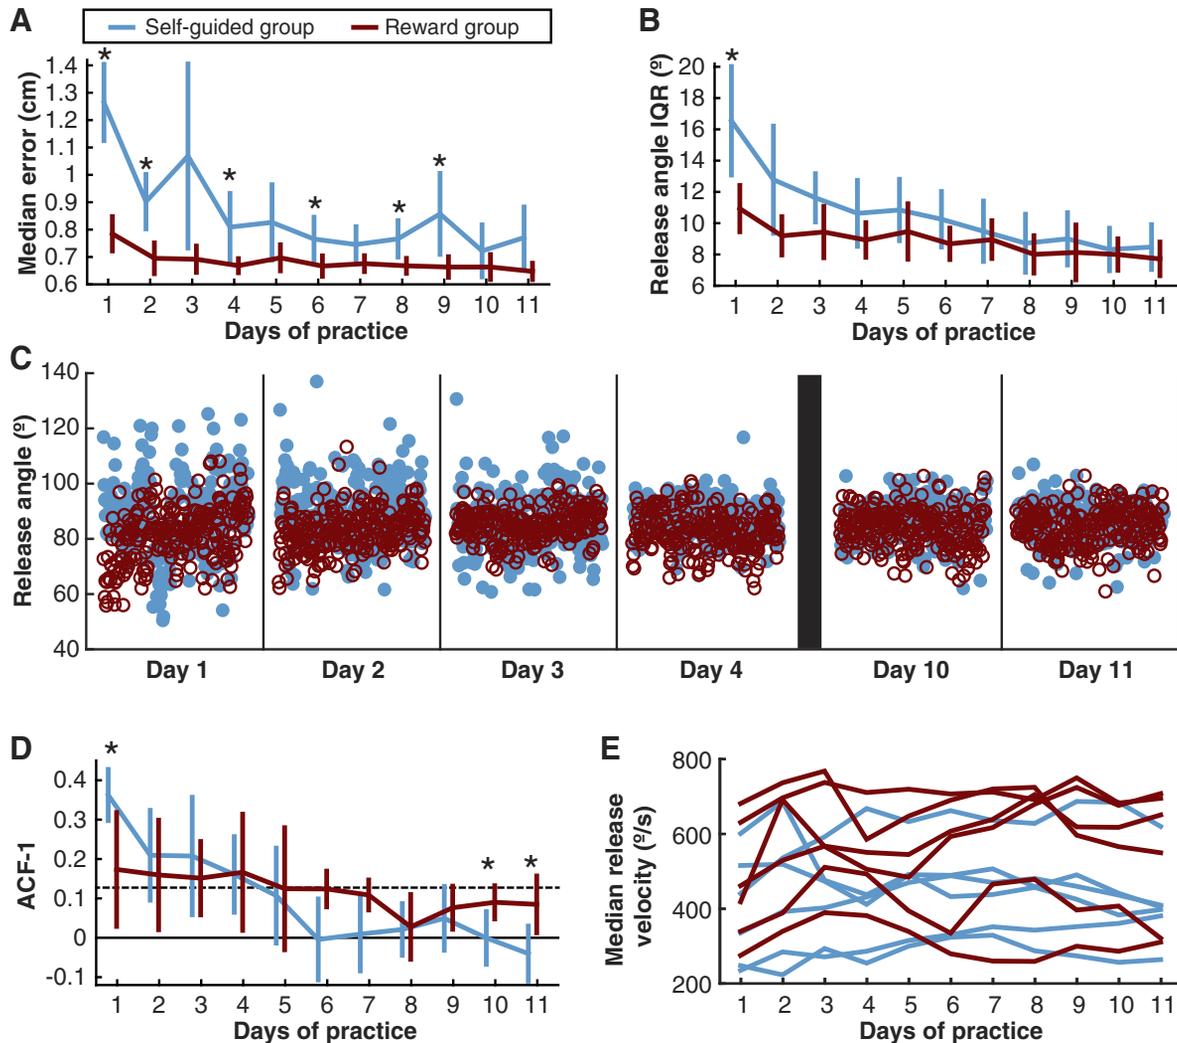


Fig. 2. *Experiment 1*: reward accelerates error reduction, but with sufficient practice, self-guided practice reaches a similar level of performance and noise. *A*: group comparison of error over 11 days of practice. The self-guided practice group had higher error for most of the practice days; however, they converged to a similar error level as the reward group by the end of practice. *B*: comparison of variability between self-guided and reward groups. While the self-guided group had higher variability on the first day of practice, this difference disappeared for the remaining 10 days of practice. *C*: time series of release angles during the first 4 days and the last 2 days of practice of an example subject in the self-guided group (blue) and the reward group (red). *D*: after 4 days of practice, the mean lag-1 autocorrelation values (ACF-1) for both groups reveal no temporal structure in release angle. This indicates that the variability in release angle after extended practice was dominated by random noise. Dashed line represents the upper 95% confidence interval for a white noise process. *E*: individual subjects did not decrease their release velocity as a strategy to decrease release angle variability. High intersubject variability results from the fact that release velocity played very little role in determining task performance. All error bars represent ± 2 SE. Significant pairwise differences between groups: $*P < 0.05$.

Self-Guided and Reward Learning Have Similar Variability

To quantify the precision of performance, the next focus of analysis was on variability of throws within each practice day. As expected, all subjects significantly reduced their variability, as measured by the IQR of release angles, over days [$F_{(10,100)} = 11.86$, $P < 0.001$, partial $\eta^2 = 0.54$; Fig. 2*B*]. However, the two groups did not differ significantly from each other [$F_{(1,10)} = 3.89$, $P = 0.077$, partial $\eta^2 = 0.28$]: while subjects in the self-guided group started with significantly greater variability on day 1 ($P = 0.014$), this group difference disappeared over the remaining 10 days (all $P > 0.075$; Fig. 2*B*). The Group \times Day interaction was nonsignificant [$F_{(10,100)} = 2.73$, $P = 0.070$, partial $\eta^2 = 0.21$]. The group differences on day 10 and day 11 were not statistically significant [day 10: $t(10) = 0.38$, $P = 0.71$, $d = 0.22$; day 11: $t(10) = 0.85$, $P = 0.42$, $d = 0.49$]. To detect the effect size observed on day 11, 67 subjects per

group would be needed, which suggests that group equivalence is not likely due to small sample size. This converging performance of the two groups is illustrated in Fig. 2*C*, showing single-trial data of release angles of an example subject in the self-guided group (blue) and reward group (red) during the first 4 days and the last 2 days of practice. These results ran counter to *hypothesis 1*.

Variability in Both Self-Guided and Reward Learning Is Random

Decreases in error and variability over practice are frequently ascribed to error corrections, particularly in motor adaptation paradigms, where error-based learning predominates. However, random fluctuations resulting from intrinsic neuromotor noise also affect these performance measures (Abe and Sternad 2013; van Beers et al. 2013). *Hypothesis 2* stated

that to improve task performance over long-term practice, noise variance has to decrease.

Initially, the average ACF-1 calculated for each daily session was positive in all subjects of both groups, indicating weak persistence across successive throws (*day 1*: mean = 0.26, SD = 0.16). Over sessions, however, ACF-1 significantly decreased and approached 0 [$F_{(10,100)} = 5.30$, $P = 0.001$, partial $\eta^2 = 0.35$; Fig. 2D]. From *day 5* onward, the average ACF-1 of all subjects and both groups was within the 95% confidence interval of uncorrelated white noise. Neither the group main effect [$F_{(1,10)} = 0.69$, $P = 0.43$, partial $\eta^2 = 0.065$] nor the interaction [$F_{(10,100)} = 1.82$, $P = 0.14$, partial $\eta^2 = 0.15$] was significant. To further test *hypothesis 2*, DFA was applied over the same data. The scaling index similarly approached 0.5, which is the value for unstructured noise, corroborating the findings of the autocorrelation results (see Fig. 3, A and B). These analyses indicate that by the end of

practice neuromotor noise predominated, consistent with *hypothesis 2*.

Reduced Variability and Noise Were Not Result of Reduced Movement Speed

While the lack of temporal structure supported the expectation that neuromotor noise attenuated with practice, alternative explanations for reducing overt variability have been proposed. For instance, a common strategy in motor performance is to lower movement speed in return for accuracy (Fitts 1954). Examining the change in the velocity at which participants released the ball did not reveal significant decrements in median release velocity across practice days [$F_{(1,100)} = 1.73$, $P = 0.19$, partial $\eta^2 = 0.15$], group differences [$F_{(1,10)} = 2.06$, $P = 0.18$, partial $\eta^2 = 0.17$], or an interaction [$F_{(1,100)} = 0.73$, $P = 0.52$, partial $\eta^2 = 0.068$] (Fig. 2E). Hence, there was no sign that a speed-accuracy trade-off could have been responsible for the observed improvements in accuracy, counter to observations in other studies (Reis et al. 2009; Shmuelof et al. 2012).

A second widely accepted mechanism underlying overt variability is that noise is dependent on signal strength (Harris and Wolpert 1998). This has been illustrated in an optimal control model and experimentally shown in isometric force production, where the increased variability at higher forces has been ascribed to signal-dependent noise (Jones et al. 2002). In dynamic tasks, high velocity requires fast muscle contractions; hence, velocity is the dynamic correlate of high force levels and faster movements are expected to be associated with higher noise. The fact that release velocity did not change with practice therefore also ruled out that subjects exploited this possibility for noise reduction (Hasson et al. 2016; Sternad et al. 2011). This lack of support for alternative explanations indirectly corroborates that the reduced variability was the result of reduced noise.

Neuromotor Noise Can Be Reduced Further with Changing Reward

In *experiment 1* reward led to faster reduction of error, but not to lower error or variability by the end of practice, counter to *hypothesis 1*. However, in support of *hypothesis 2*, noise processes dominated the observed variability in the sequence of throws by the end of practice. This second finding may imply that subjects had reached their physiological limit. Alternatively, subjects may only have reached a point of diminishing returns, where any further reduction of noise was not worth the effort, especially given their high success rate (84% of trials were successful by the end of practice). To discern between these two possible explanations, *experiment 2* manipulated the reward threshold to make it harder to achieve reward. *Hypothesis 3* posited that the higher demands would elicit further decreases in the amplitude of noise. As a direct test, *experiment 2* first presented the same threshold for reward as in *experiment 1* over *days 1–3* but changed the threshold over *days 4–6* to require more accurate throws to receive reward (Fig. 4A). We expected that the noise amplitude would be lowered.

To further test that practice induced persistent effects, the tighter task demands were relaxed again over *days 7–11*. While in motor adaptation the acquired behavior invariably returns to

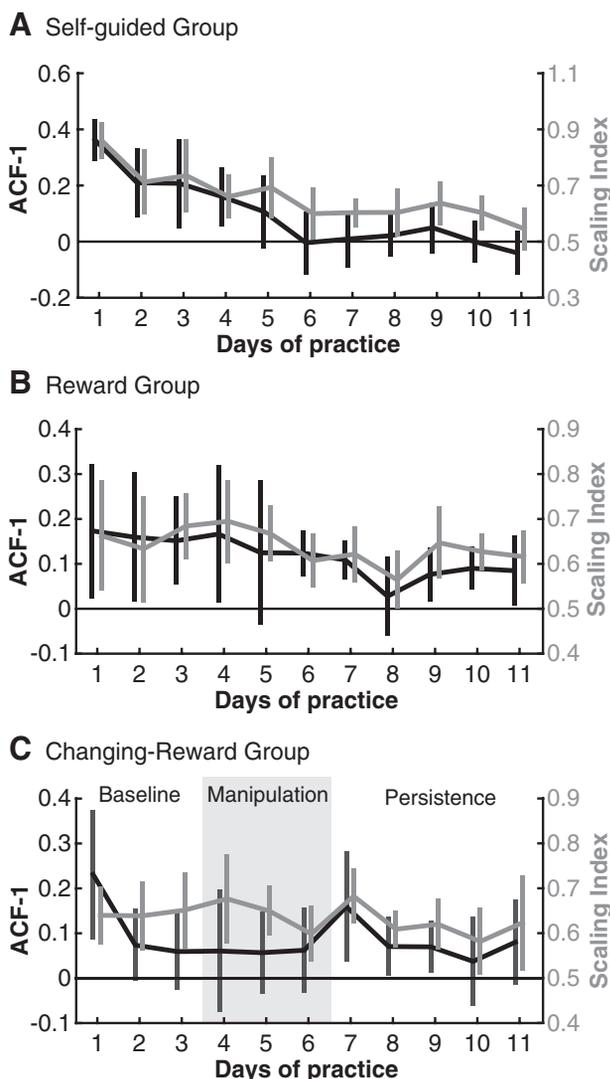


Fig. 3. *Experiment 1*: autocorrelation and detrended fluctuation analyses consistently show that trial-to-trial fluctuations converged to a noise process. A–C: autocorrelation and detrended fluctuation analyses were conducted for each experimental group. ACF-1 values close to 0 and scaling exponent values of 0.5 indicate unstructured, random noise. Both analyses produced similar results. All error bars represent ± 2 SE.

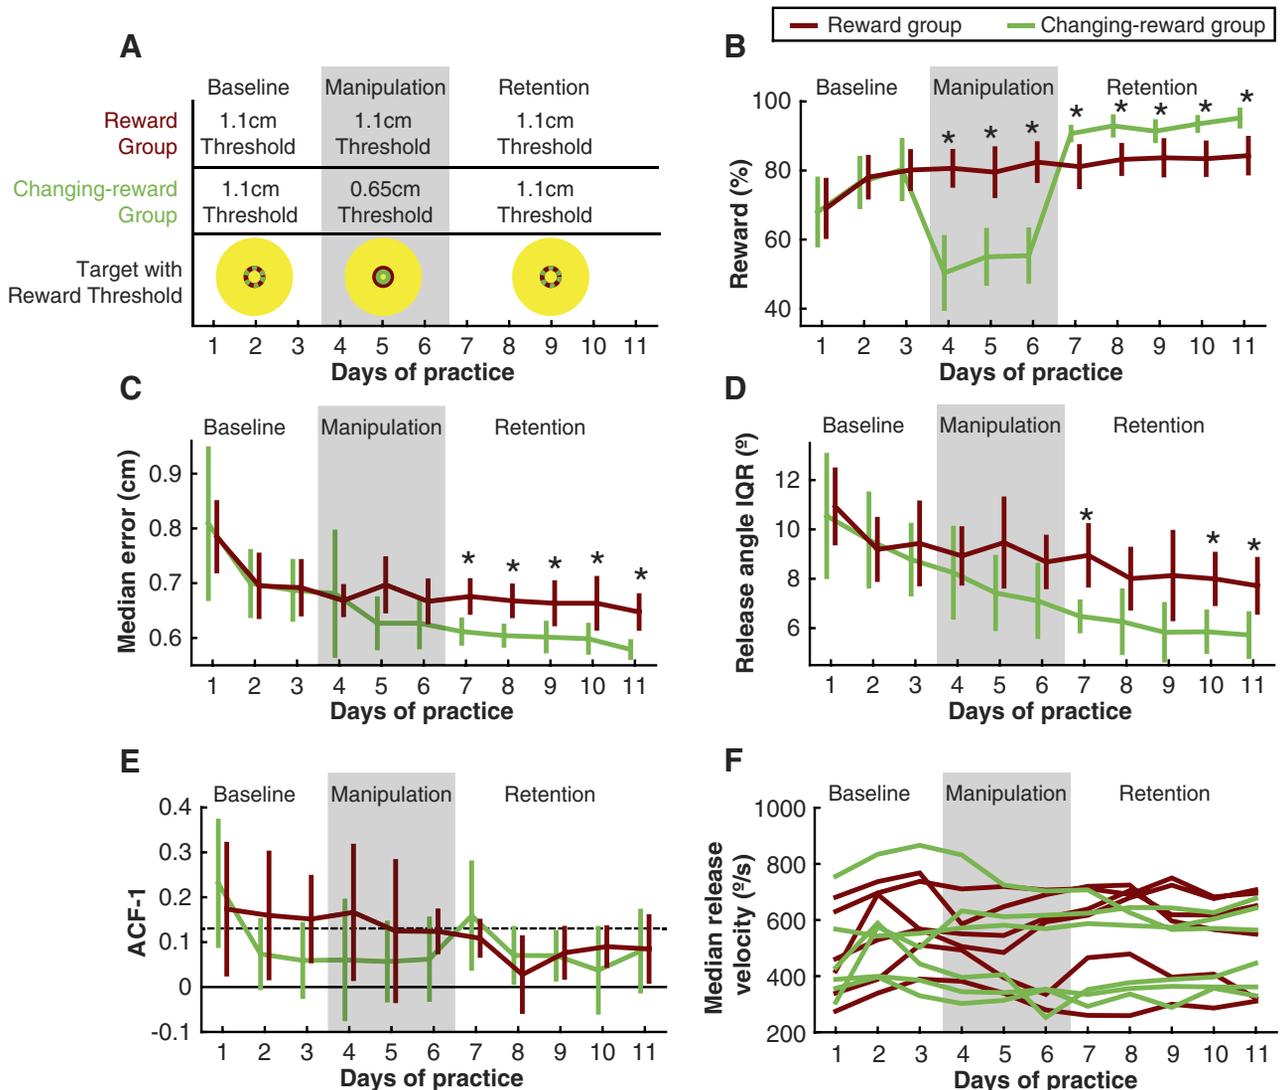


Fig. 4. *Experiment 2*: the reward condition that requires more accurate performance leads to lower variability that persists over 5 additional practice days. *A*: the reward group, which is the same as in *experiment 1*, practiced with a reward threshold of 1.1 cm for all 11 days of practice. The changing-reward group practiced with a reward threshold of 1.1 cm for the first 3 days of practice, followed by 3 days with a more challenging reward threshold of 0.65 cm and 5 days with the original threshold of 1.1 cm. *B–E*: group comparisons of dependent measures. The shaded region indicates the practice days where the reward threshold between the 2 groups differed. *B*: % of rewarded throws on each practice day was reduced with the more challenging threshold as expected for the changing-reward group. *C* and *D*: while group differences in error and variability were not significantly different during the reduced reward threshold as expected, they did emerge during the persistence period, as the changing-reward group performed the task with lower error and lower variability. *E*: average ACF-1 of all subjects was within the 95% confidence interval of uncorrelated white noise (dashed line) from the second day of practice onward. *F*: as in *experiment 1*, individual subjects did not decrease their release velocity as a strategy to decrease release angle variability. All error bars represent ± 2 SE. Significance: $*P < 0.05$.

the original behavior after removal of the external perturbation, the hallmark of skill learning is that the learned behavior persists for a long time, indicative of neuroplasticity. Hence, to test that practice induced persistent effects, the initial threshold was reinstated on *day 7* and practiced for 5 more days until *day 11* (Fig. 4A). *Hypothesis 4* stated that subjects would retain their acquired low level of noise.

On *days 1–3*, subjects in both groups significantly improved task performance, marked by an increasing percentage of successful trials [$F_{(2,20)} = 17.70$, $P < 0.001$, partial $\eta^2 = 0.64$; Fig. 4B]. Corroborating results of *experiment 1*, both groups decreased median error [$F_{(2,20)} = 10.04$, $P = 0.003$, partial $\eta^2 = 0.50$; Fig. 4C] and release angle variability [$F_{(2,20)} = 4.85$, $P = 0.039$, partial $\eta^2 = 0.33$; Fig. 4D] without any

difference between them. The temporal structure in release angle variability as measured by ACF-1 did not show any significant change across days [$F_{(2,20)} = 2.76$, $P = 0.099$, partial $\eta^2 = 0.22$; Fig. 4E]. The average ACF-1 of all subjects was within the 95% confidence interval of uncorrelated white noise on *days 2* and *3*, suggesting that variability was due to noise on these days. The autocorrelation findings were corroborated by the results of the DFA (see Fig. 3, B and C). As expected, there were no significant effects of group or Group \times Day interaction on any of the dependent measures (all $F < 1.8$), repeating the results of *experiment 1*. Although variability in error and release angle decreased, release velocity did not change as a function of practice as in *experiment 1*. In fact, the median release velocity slightly increased over the 3 days

$[F_{(2,20)} = 9.66, P = 0.003, \text{partial } \eta^2 = 0.49; \text{ Fig. } 4F]$. This suggested that variability did not decline as a result of slower movement speeds as expected from Fitts's law.

During *days 4–6*, subjects in the changing-reward group were exposed to an altered threshold of 0.65 cm to enhance the task challenge; the threshold remained at 1.1 cm for the reward group. Again, the visual size of the target was unchanged and subjects were not informed about this change (the rings in Fig. 4A only illustrate the threshold for the reader). As expected, the success rate of the changing-reward group dropped significantly compared with the reward group [$F_{(1,10)} = 31.92, P < 0.001, \text{partial } \eta^2 = 0.76; \text{ Fig. } 4B$], confirming that task success was indeed more difficult for these subjects. Median error and release angle variability showed trends suggesting that the changing-reward group improved faster than the reward group as hypothesized, although these interactions did not reach statistical significance (all $P < 0.17$). There were no significant effects or interactions in any other dependent measures (all $P > 0.05$).

Persistence of Reduced Noise

While subjects only gradually responded to the change in the reward threshold, a significant group difference did emerge on *day 7* when the threshold was restored to the original value. Even though both groups practiced with the larger reward threshold for the remaining 5 days, the changing-reward group continued with a significantly higher success rate [$F_{(1,10)} = 19.06, P = 0.001, \text{partial } \eta^2 = 0.66; \text{ Fig. } 4B$]. This improved performance resulted from significantly lower error [$F_{(1,10)} = 19.97, P = 0.001, \text{partial } \eta^2 = 0.66; \text{ Fig. } 4C$] and variability [$F_{(1,10)} = 10.02, P = 0.01, \text{partial } \eta^2 = 0.50; \text{ Fig. } 4D$]. These results were consistent with *hypothesis 4*. The temporal structure of variability did not differ between the two groups (all $P > 0.05$), and the average ACF-1 for all subjects in both groups was within the 95% confidence interval of uncorrelated white noise for all 5 days (Fig. 4E). No significant effect of day or interaction was observed in any of the measures (all $P > 0.05$). There was no group difference in median release velocity [$F_{(1,10)} = 0.66, P = 0.44, \text{partial } \eta^2 = 0.06$], which again ruled out that speed-accuracy trade-off or lower signal strength may have lowered the overt variability or noise.

These behavioral results suggest that it was physiologically possible for subjects in the changing-reward group to decrease the noise in their motor behavior and achieve better task performance.

Simple Model with Time-Varying Gain on Noise Reproduces Empirical Results

Model with constant noise. To illustrate how the change in success requirements could influence motor variability via its two essential components, error correction and noise, we used an extremely simple yet prevalent iterative learning model. With this model, we aimed to further probe into the effects of error corrections and noise on execution variability, particularly at the learning asymptote, and how can it be reduced. The model describes how the central nervous system might correct movement errors on a trial-by-trial basis:

$$x(n+1) = x(n) - Be(n) + \xi(n) \quad (1)$$

$$e(n) = x(n) - x^* \quad (2)$$

where $x(n)$ is the motor output at trial n and $e(n)$ is the error at trial n , defined as the difference between the actual and desired motor output x^* . The parameter B defines the fraction of error corrected, or error correction gain, ξ is an additive white Gaussian noise term, $\xi \sim N(0, \sigma^2)$. This model in its present form and with slight variations has been successful in simulating a range of phenomena associated with motor adaptation (Herzfeld et al. 2014; Smith et al. 2006) and learning (van Beers et al. 2013).

On the basis of this model, how can variability be decreased? In principle, there are two options: either increase the error correction gain or decrease the variance of the noise process. To determine whether an increase in error correction gain could explain the decrease in variability, we simulated data using *Eqs. 1 and 2* with five different B values while keeping the noise variance ξ constant. The motor output $x(n)$ in the model represented the release angle for a given trial n ; x^* was the optimal release angle that resulted in zero error; $e(n)$ represented the error in release angle. For simplicity we refrained from including the mapping between release angle and error in the simulations. Although error in release angle is different from the error measure used in the experiments, these two error measures are tightly related. As shown in Fig. 1G, the error increases with the difference between the actual and the optimal release angle, although not linearly for larger errors.

For each value of B (0.1, 0.3, 0.5, 0.7, 0.9), the model was simulated for 2,640 trials to match the total number of trials performed by the human subjects. The results were analyzed in bins of 240 trials representing “days” to match the analysis of the experimental data. The optimal angle x^* for this target location was 82.44° (Fig. 1E). The average median release angle on *day 1* from subjects in the reward and changing-reward groups was used as initial value of x , such that $x(1) = 83^\circ$. The average variance of release angle on *day 1* from all subjects in the reward and changing-reward groups was used to set the variance of the additive white Gaussian noise, such that $\sigma^2 = 68.9^\circ$.

Figure 5A shows that variability in release angle is lowest for higher error correction gains ($B \sim 1$). Importantly, the magnitude of variability remained relatively constant over simulated days of practice regardless of the value of B . Additional simulations with different values of noise variance ($30^\circ < \sigma^2 < 100^\circ$) confirmed that the relation between the correction gain B and ACF-1 is insensitive to changes in the noise variance. Hence, different magnitudes of noise variances consistently produced the same pattern of results. To decrease variability through the error correction gain alone, B must increase over trials, i.e., it must become a function of trials, $B(n)$.

However, Fig. 5B also demonstrates how different B values produced different temporal structure in the variability: the higher the error correction gain, the closer to 0 the ACF-1 value becomes (Abe and Sternad 2013; van Beers 2009; van Beers et al. 2013). In *experiment 2*, the ACF-1 value was within the range of white noise from *day 2* onward, suggesting that the error correction gain B was already high (~ 0.90) at the beginning of practice. Furthermore, the ACF-1 did not significantly change over days. Thus it is unlikely that the reduced variability observed in the experimental data was achieved by an increasing error correction gain. This suggests that the noise variance had to be scaled.

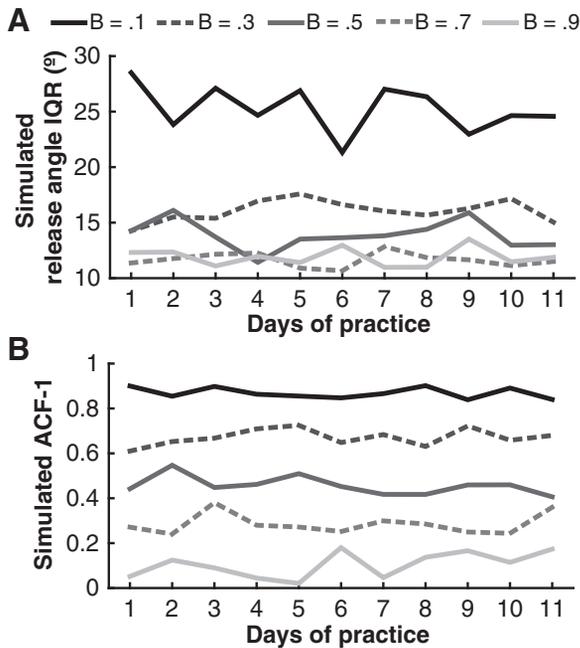


Fig. 5. *A*: simulated data from the iterative model with constant noise term (Eqs. 1 and 2) and different B values while keeping the noise amplitude constant. For each simulated value of B , variability measured by the interquartile range remained relatively constant over simulated days of practice. *B*: different B values result in different temporal structure in the variability. For higher error correction gains, ACF-1 approaches 0.

Model with constant time-varying noise gain. To explain the results of *experiment 2*, we introduced a gain factor $a(n)$ to modulate the magnitude of the noise variance as follows:

$$x(n+1) = x(n) - Be(n) + a(n)\xi(n) \quad (3)$$

Next, we defined how this noise gain changed over practice. For both reward groups, the results of *experiment 2* showed a clear relation between the reward threshold and the magnitude of random variability: subjects attenuated noise to increase reward. This finding concurred with previous results in a motor adaptation task by Izawa and Shadmehr (2011), who demonstrated that reward could be used to modulate noise. However, in their study, subjects amplified noise to explore solutions that yielded a higher reward. In the present study, subjects were already centered on the optimal solution (Fig. 1*E*). To exploit this correct solution and maximize reward, it was advantageous to decrease noise.

Therefore, the following rule for updating the noise scaling factor $a(n)$ was added:

$$a(n+1) = (1 - 1_{e(n) > T}A)a(n) \quad (4)$$

where $e(n) > T$ denotes the condition that error is greater than the reward threshold, T . The term $1_{e(n) > T}$ denotes an indicator function that equals 1 if the absolute value of error $e(n)$ is greater than the threshold T and no reward is given. The function is set to 0 if reward is received. A denotes the rate at which the noise scaling factor decreases. According to this update rule, $a(n)$ is decreased whenever no reward is obtained. Figure 6 shows results from this model (Eqs. 2–4) together with the results of *experiment 2*. Note that in the experiment the reward threshold for success was based on error. For the simulation, the error threshold was transformed into a corre-

sponding angle threshold. A 1.1-cm error threshold corresponded to an 8.9° angle threshold, and a 0.65-cm error threshold corresponded to a 4.4° angle threshold. The value of B was set to 0.86 and A was set to 0.001 to match the temporal and spatial characteristics of the variability observed in the experimental data (see *Model fitting and parameter sensitivity*), and $a(1)$ was set to 1. The model replicated the pattern of the reward percentage in both the reward and changing-reward groups (Fig. 6*A*). It also replicated the slow decline in release angle variability after the change in threshold on *days 4–6* in the changing-reward group. Importantly, it shows that the release angle variability was maintained after the initial reward threshold was reinstated (Fig. 6*B*). ACF-1 showed negligible variations across the practice days in both experimental and simulated data (Fig. 6*C*).

Note that the self-guided group also scaled down the noise over practice, eventually reaching the level of the constant-reward group. However, it is unclear from the experimental protocol what conditions were responsible for the decrease of $a(n)$ from one trial to the next. While self-guided learning is evidently ubiquitous in the real world, studying this “uncon-

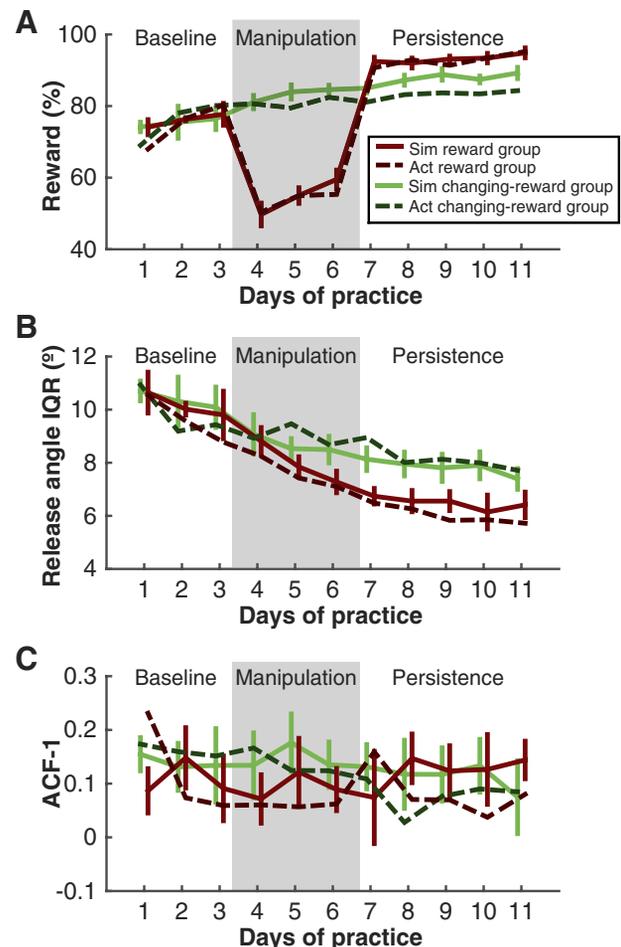


Fig. 6. Simulation of decreased neuromotor noise as a function of binary reward reproduces empirical results. Incorporating the update rule that decreases the noise amplitude as a function of reward (Eqs. 2–4) reproduced the empirical results of % of rewarded trials (*A*), release angle IQR (*B*), and ACF-1 (*C*) in *experiment 2*. Six simulations were conducted for each group. All error bars represent ± 2 SE.

trolled" form of learning with a model-based approach remains a challenge.

Model fitting and parameter sensitivity. In this model, two parameters had to be fit to the experimental data: the error correction gain, B , and the rate at which the noise scaling factor changes, A . To fit these parameters, the following procedure was implemented: for given values of B and A , six simulations were conducted for each group. For each group, the average release angle IQR and ACF-1 of release angles were calculated on each day (i.e., bins of 240 throws). For each measure, the squared differences between the simulated and actual group averages on each practice day were summed. As the goal was to fit the model output to both experimental groups, the sum of the squared differences for each group were added together. This measure served as the fit metric: the lower the fit metric, the better the simulated data matched the experimental data of both groups. This procedure was repeated to generate one fit metric for release angle IQR and another for ACF-1 for all combinations of A between 0 and 0.1 and B between 0 and 1.

Figure 7, A and B , illustrate how the different values of A and B affect the fit of the simulated data to the experimental data. The darker values depict the combination of model parameters with the best, i.e., lowest fit metric, whereas the lighter values depict the combinations with the worst fits. When determining the model parameters that best fit the experimental data, it was important to take into account how well the model output reproduced the trends of both release angle IQR and ACF-1 in the experimental data. Because the units of the fit metric depended on the units of the dependent variable, these two metrics could not simply be added. Instead, each combination of model parameters was given a ranking for each of the two dependent measures based on their fit metric. These rankings were summed, and the combination of A and B values that resulted in the lowest summed rank were chosen as the parameters that best fit the experimental data. The best values were $A = 0.001$ and $B = 0.86$ (Fig. 7C).

The results of this fitting procedure also demonstrated one of the most important features of this model with noise modulation: the spatial and temporal characteristics of the variability in model output could be independently controlled. For values of B between 0.32 and 1, the A parameter solely determined how well the simulated data fit the experimental data in terms of release angle IQR (Fig. 7A). While the value of the B parameter had little influence on release angle IQR within this range, it did profoundly affect the ACF-1 of the model output (Fig. 7B). In contrast, the value of the A parameter did not affect the fit of ACF-1. Changing A only affects the magnitude of the noise variance; it does not change its temporal structure. As ACF-1 is a measure of the average relation between two consecutive trials, it is not influenced by the magnitude of the additive noise. These results demonstrate how the two free parameters of the model could be independently tuned to match spatial (parameter A) and temporal (parameter B) characteristics of the experimental data.

DISCUSSION

Mastering a motor skill can take a lifetime. And yet, most motor learning and adaptation studies have examined only the initial phases of learning. In contrast, this study focused on the later fine-tuning processes of skill learning, where only subtle

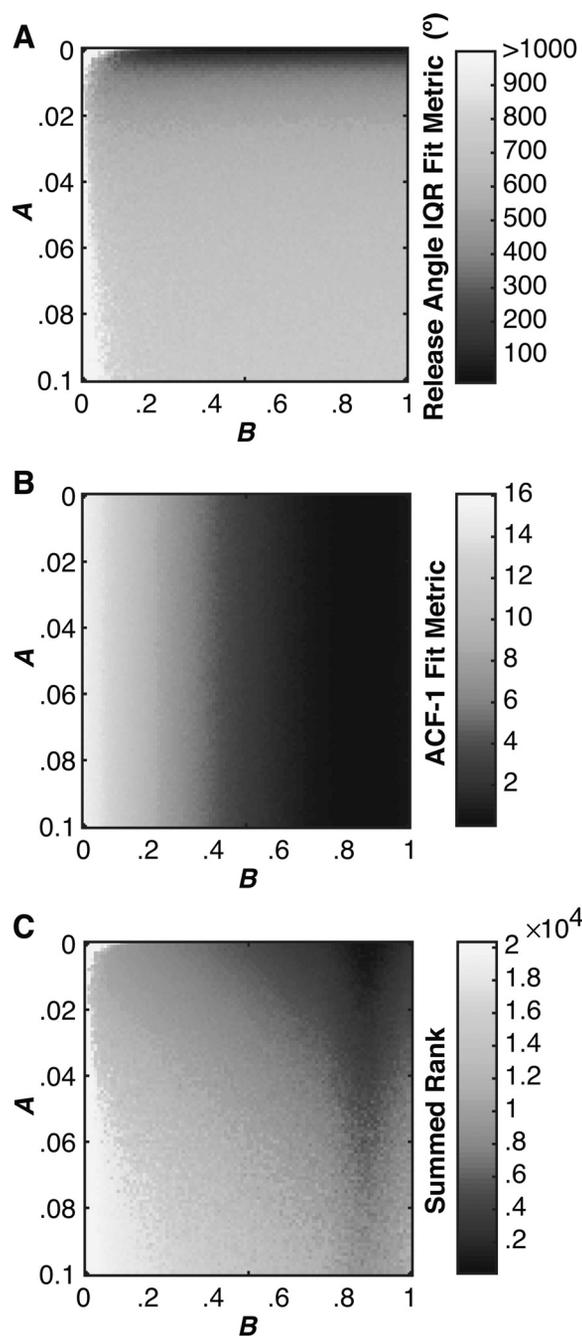


Fig. 7. Model fitting and parameter sensitivity analysis. A and B : for each combination of the 2 model parameters B and A , the fit metric for release angle IQR (A) and ACF-1 (B) is depicted by color (darker shades indicate better fits). C : for each combination of the 2 model parameters B and A , the ranks of the 2 corresponding fit metrics for release angle IQR and ACF-1 were summed. This summed rank is depicted by color (darker shades indicate lower ranks). The combination that generated the lowest rank value ($A = 0.001$ and $B = 0.86$) produced the best fit for experimental data.

changes occur during many hours of practice. The overall hypothesis was that the characteristic decline in overt variability with practice is due to an attenuation of the magnitude of neuromotor noise rather than increased error corrections. To enhance this slow tuning process, subjects were given reward for additional feedback and motivation. We further hypothesized that with reward not only would the magnitude of these noise processes decrease, but this attenuation would also per-

sist for an extended time, indicative of lasting plastic changes in the nervous system. We assumed that overt variability is generated by feedback-based error corrections and by neuromotor noise. At later stages of practice, analysis of our behavioral data found no evidence for increased error corrections. Hence, we rejected the alternative hypothesis that error corrections gave rise to the measured variability.

Previous studies by our group have already examined long-term practice with the same skittles task to monitor changes in variability over 6 and even 16 days of practice (Abe and Sternad 2013; Cohen and Sternad 2009; Müller and Sternad 2004). The goal of these earlier studies was to decompose variability in sets of data, rather than assessing fluctuations in their temporal sequence. Parsing the distributional pattern of sets of throws in the execution space, analyses identified three components of performance improvement: Tolerance, Covariation, and Noise (Müller and Sternad 2009; Sternad and Abe 2010). Results showed that with practice subjects first optimized Tolerance, i.e., identified solutions where error and noise affected performance least. Over longer practice, they aligned their variability with the solution manifold (Covariation), again to make errors and noise matter less. This stage in the learning process is consistent with the minimum intervention principle, exemplified in models using stochastic optimal feedback control (Todorov and Jordan 2002).

Important for this study, the contribution of the Noise component decreased extremely slowly and remained the highest even after long practice. This finding motivated the present study to investigate whether additional interventions, such as giving reward, could reduce the Noise component. The present study deliberately modified the target constellation to create a solution manifold where release angle and velocity could not covary, i.e., Covariation was eliminated as a way to attenuate the effect of variability or noise on the result. This solution manifold also eased the challenge and made Tolerance less important, which allowed full focus on the random noise component.

Frequent Reward Accelerates Learning But Does Not Lead to Better Overall Performance

According to many motor learning textbooks, augmented feedback is arguably the single most important variable for motor learning (Salmoni et al. 1984; Schmidt and Lee 2011). However, the experimental results qualified this straightforward expectation and highlighted that while augmented feedback did promote learning in the early stage of practice, in accord with *hypothesis 1*, it failed to provide an advantage over self-guided practice in the longer run. While these results may benefit from more statistical power, the finding resonates with numerous examples in real-world scenarios and experimental settings in previous studies, where learning proceeds without controlled external feedback and reward (Park et al. 2013; Park and Sternad 2015). We speculate that internal feedback mechanisms via visual and proprioceptive information proved equally effective (Cordo 1990).

While the results for error were more differentiated, reward provided little benefit for the IQR of release angle after initial practice. This result is noteworthy, as studies on reward have paid little attention to its influence on variability compared with average behavior. In adaptation, for example, reward or

reinforcement learning has been used to guide a subject's behavior toward a new solution. Hence, measures of central tendency (e.g., average signed error from the task goal) have been more frequently assessed than measures of dispersion, although recently Nikooyan and Ahmed (2015) reported that the combination of reward and sensory feedback can accelerate the reduction of variability during reaching adaptation. It should be noted, however, that practice in that study was limited to one practice session, as is typical for most adaptation studies. Our results similarly suggest that frequent reward lowers variability over the course of one practice day, but its effect diminishes over long-term practice.

Our results further showed that when the reward became sufficiently challenging and less frequent, it could provide an advantage over self-guided practice. Apparently, subjects had not reached their limit, as increased accuracy requirements for the success signal could elicit further improvements (*hypothesis 3*). One reason may be that the cost to lower variability or noise may have been too high, consistent with the assertion of Manohar et al. (2015). Alternatively, there may also be advantages to variability in motor behavior, such as supporting exploration. Recently, Wu et al. (2014) provided renewed evidence that initial variability can be beneficial for exploring the solution space to find the behavior that receives reward, although this effect may be complex and task specific (He et al. 2016). Yet, even after the initial exploration, it is possible that subjects maintain a certain level of variability in order to gain information about the task (Kaelbling et al. 1996). Frequent reward may have discouraged exploration and instead prioritized exploitation. It has been also hypothesized that it is not the reward per se that drives learning but rather the mismatch between the expectation and receipt of a reward (Tobler et al. 2006). The fact that the changing-reward group did not receive reward for trials that were previously rewarded may explain why this more difficult reward led subjects to further reduce their variability.

Neural Mechanisms to Reduce Neuromotor Noise

While it is not surprising that variability decreased under tighter task demands, it is nontrivial that it was the decrease in the variance of the random processes, not an increase in error correction, that reduced the overt variability (*hypotheses 2 and 3*). How might the central nervous system reduce not only overt but possibly also physiological noise? One potential mechanism is to increase coactivation in antagonistic muscles (Gribble et al. 2003; Selen et al. 2009), although antagonistic cocontraction tends to decrease with practice (Thoroughman and Shadmehr 1999). Lower signal-dependent noise has been ruled out as a candidate by the experimental results that showed no decrease in velocity (Harris and Wolpert 1998). Another conjecture is derived from studies on the effect of neuromodulators on motor neuron excitability, such as serotonin (Theiss and Heckman 2005; Wei et al. 2014), norepinephrine (Theiss and Heckman 2005), and dopamine (Kroener et al. 2009). Animal studies provided intriguing evidence that the descending drive to muscle contractions is gain controlled to modulate the required output force. One study on humans specifically showed that force variability increased after the brain stem-spinal cord neuromodulatory system was upregulated (Wei et al. 2014). The complex interplay of neuromodu-

lators can excite or inhibit spinal cord excitability and thereby may match precision demands in motor tasks. Finally, Picard and colleagues demonstrated in skilled performance of monkeys that years of practice resulted in more efficient generation of neuronal activity in M1 (Picard et al. 2013). It is feasible that such a decrease in metabolic activity could also result in lowered neuromotor noise. Alternatively, the variability in neuronal activity may be shaped and reduced in neural space and in a way that leads to more consistent performance (Sadler et al. 2014). Evidently, more research is needed to solidify these conjectures.

Long-Term Persistence Indicative of Long-Term Neuroplasticity

Experiment 2 tested not only whether modulation of reward threshold influenced learning but also whether this effect showed persistence over the 5 days after the reward criterion was relaxed. While *hypothesis 4* stated that random processes would remain suppressed, the experimental result was nevertheless stronger than expected. At first blush, it appeared similarly reasonable that random fluctuations might have increased after the task constraints were alleviated, as previous studies provided ample documentation that humans shape their variability rather than only reducing it (Chu et al. 2013; Cohen and Sternad 2009; Müller and Sternad 2004; Sternad et al. 2014). However, having suppressed the magnitude of random fluctuations, subjects may not have sensed that relaxing their performance variability was possible, as there was no verbal or visual indication. Alternatively, and following *hypothesis 4*, subjects may indeed have permanently altered their performance, because of lasting changes in the nervous system.

Note that this persistence differs from savings, which is one sign of neuroplastic effects in motor adaptation; savings refers to the reduced time for readapting when exposed to the same perturbation more than once. In adaptation, subjects experience large errors when the perturbation is removed and, consequently, the original behavior is reinstated to perform the task successfully. In the present study, the acquired behavior was the correct behavior in all conditions. Thus, the learned behavior with reduced variability could persist without penalty after the manipulation was removed. This result underscores that interventions to enhance long-term behavior should use conditions that enhance a behavior that remains essentially identical, even when the intervention is removed (Huber and Sternad 2015; Reisman et al. 2009).

Considerations and Limitations of Modeling Skill Learning

To illustrate the interplay of error, correction gain, and noise in trial-to-trial learning, an extremely simple recursive model generated results that matched the behavioral data observed in *experiment 2*. The model combined the usual error correction gain with an additive noise term and a time-varying noise gain that was altered by a simple indicator function representing the threshold for reward. The model successfully simulated the decrease in variability and the persistence of the low level of noise. Interestingly, it also captured the delayed response to the changed reward that only became significant after 3 days of practice. Note that the present model included only a single additive noise source and no multiplicative noise source.

Evidently, modeling results rely heavily on the structure and order of the model. In recent studies on motor learning, most models have focused on the deterministic processes, i.e., error correction (Herzfeld et al. 2014; Smith et al. 2006). Several studies included stochastic iterative models with one or two noise terms. Specifically, van Beers (2009) and Abe and Sternad (2013) showed that at least two noise sources were required to account for the observed temporal structure in the data. A recent study by Hasson et al. (2016) examined the effect of error amplification and used system identification techniques with three different iterative models. Estimating the error correction gain B and the noise sources revealed that the primary change in the observed error was accounted for by a decrease in the variance of the noise sources. This result is consistent across all three models, including the simplest model with one noise source and the two models with two noise sources (van Beers 2009). The difference between these models was the change in the correction gain, which showed only a slight decrease over the extended practice, consistent with a previous study on skittles learning (Abe and Sternad 2013).

Despite the simplicity of the model used in the present study, the simulation demonstrated how variability and particularly the noise from neuromotor system might decrease in response to the change of reward feedback. The behavioral results of the self-guided group also showed a decrease in variability over practice. However, what drove these subjects to lower their variability remains unknown, as there was no comparison group to test a hypothesized mechanism. One challenge with modeling self-guided, or unsupervised, skill learning is that, by definition, there is no explicit or added performance feedback given during learning (Wolpert et al. 2001). Thus, it is difficult to identify, and consequently simulate, the exact information subjects are using to update their performance from one trial to the next. Nevertheless, efforts to better understand this type of skill learning are needed because this is probably the most prevalent form of practice in the real world. As with any modeling attempt, behavioral and physiological data are needed to reveal the patterns of behavior when we acquire and retain complex motor skills.

Modeling long-term skill learning that addresses the ever-present stochastic components still remains an open challenge in the field of motor neuroscience. Neural networks present an alternative to explore the role of noise in skill acquisition, long-term retention, and generalization (Ajemian et al. 2013). However, stochastic models and system identification are still in their infant stages in motor neuroscience, despite the recognized critical role of noise in the nervous system.

Conclusions

Long-term skill learning and retention requires attention as it opens an interesting view on a ubiquitous but ill-understood element of the neuromotor system—noise. Our experimental results suggest that skill learning proceeds not only by correcting errors but also by reducing the amplitude of neuromotor noise, especially at later stages of practice. Reward enhanced this change, but self-guided practice was also effective. The intriguing finding was that this enhanced performance persisted for days after the reward threshold was increased. This study highlights the need to further understand neural and computa-

tional processes underlying long-term learning, especially as we continue to look for avenues to enhance skill learning and expedite rehabilitation.

GRANTS

This work was supported by National Institute of Child Health and Human Development Grants R01-HD-045639, R01-HD-081346, and R01-HD-087089 and National Science Foundation NSF-DMS 0928587 and NSF-EAGER 1548514, awarded to D. Sternad. Fellowships from the Northeastern University Graduate School of Engineering and The MathWorks supported M. E. Huber.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

AUTHOR CONTRIBUTIONS

M.E.H. and D.S. conceived and designed research; M.E.H. performed experiments; M.E.H. and N.K. analyzed data; M.E.H., N.K., and D.S. interpreted results of experiments; M.E.H. prepared figures; M.E.H. drafted manuscript; M.E.H., and D.S. edited and revised manuscript; M.E.H., N.K., and D.S. approved final version of manuscript.

REFERENCES

- Abe M, Schambra H, Wassermann EM, Luckenbaugh D, Schweighofer N, Cohen LG. Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr Biol* 21: 557–562, 2011.
- Abe MO, Sternad D. Directionality in distribution and temporal structure of variability in skill acquisition. *Front Hum Neurosci* 7: 225, 2013.
- Ajemian R, D'Ausilio A, Moorman H, Bizzi E. A theory for how sensorimotor skills are learned and retained in noisy and nonstationary neural circuits. *Proc Natl Acad Sci USA* 110: E5078–E5087, 2013.
- van Beers RJ. Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63: 406–417, 2009.
- van Beers RJ, van der Meer Y, Veerman RM. What autocorrelation tells us about motor variability: insights from dart throwing. *PLoS One* 8: e64332, 2013.
- Chu VW, Sternad D, Sanger TD. Healthy and dystonic children compensate for changes in motor variability. *J Neurophysiol* 109: 2169–2178, 2013.
- Cohen RG, Sternad D. Variability in motor learning: relocating, channeling and reducing noise. *Exp Brain Res* 193: 69–83, 2009.
- Cohen RG, Sternad D. State space analysis of timing: exploiting task redundancy to reduce sensitivity to timing. *J Neurophysiol* 107: 618–627, 2012.
- Cordo PJ. Kinesthetic control of a multijoint movement sequence. *J Neurophysiol* 63: 161–172, 1990.
- Crossman ER. A theory of the acquisition of speed-skill. *Ergonomics* 2: 153–166, 1959.
- Dingwell JB, Cusumano JP. Re-interpreting detrended fluctuation analyses of stride-to-stride variability in human walking. *Gait Posture* 32: 348–353, 2010.
- Faisal AA, Selen LP, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008.
- Fitts PM. The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47: 381–391, 1954.
- Galea JM, Mallia E, Rothwell J, Diedrichsen J. The dissociable effects of punishment and reward on motor learning. *Nat Neurosci* 18: 597–602, 2015.
- Gribble PL, Mullin LI, Cothros N, Mattar A. Role of cocontraction in arm movement accuracy. *J Neurophysiol* 89: 2396–2405, 2003.
- Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Hasson CJ, Zhang Z, Abe MO, Sternad D. Neuromotor noise is malleable by amplifying perceived errors. *PLoS Comput Biol* 12: e1005044, 2016.
- He K, Liang Y, Abdollahi F, Bittmann MF, Kording K, Wei K. The statistical determinants of the speed of motor learning. *PLoS Comput Biol* 12: e1005023, 2016.
- Herzfeld DJ, Vaswani PA, Marko MK, Shadmehr R. A memory of errors in sensorimotor learning. *Science* 345: 1349–1353, 2014.
- Hore J, Watts S, Martin J, Miller B. Timing of finger opening and ball release in fast and accurate overarm throws. *Exp Brain Res* 103: 277–286, 1995.
- Howell DC. Median absolute deviation. In: *Wiley StatsRef: Statistics Reference Online*, 2014.
- Huber ME, Sternad D. Implicit guidance to stable performance in a rhythmic perceptual-motor skill. *Exp Brain Res* 233: 1783–1799, 2015.
- Izawa J, Shadmehr R. Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput Biol* 7: e1002012, 2011.
- Jones KE, Hamilton AF, Wolpert DM. Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88: 1533–1544, 2002.
- Jørgensen HS, Nakayama H, Raaschou HO, Vive-Larsen J, Støier M, Olsen TS. Outcome and time course of recovery in stroke. II. Time course of recovery. The Copenhagen Stroke Study. *Arch Phys Med Rehabil* 76: 406–412, 1995.
- Kaelbling LP, Littman ML, Moore AW. Reinforcement learning: a survey. *J Artif Intell Res* 4: 237–285, 1996.
- Kirk R. *Experimental Design*. Pacific Grove, CA: Brooks/Cole, 1995.
- Kroener S, Chandler LJ, Phillips PE, Seamans JK. Dopamine modulates persistent synaptic activity and enhances the signal-to-noise ratio in the prefrontal cortex. *PLoS One* 4: e6507, 2009.
- Manohar SG, Chong TT, Apps MA, Batla A, Stamelou M, Jarman PR, Bhatia KP, Husain M. Reward pays the cost of noise reduction in motor and cognitive control. *Curr Biol* 25: 1707–1716, 2015.
- Müller H, Sternad D. Decomposition of variability in the execution of goal-oriented tasks: three components of skill improvement. *J Exp Psychol Hum Percept Perform* 30: 212–233, 2004.
- Müller H, Sternad D. Motor learning: changes in the structure of variability in a redundant task. *Adv Exp Med Biol* 629: 439–456, 2009.
- Nikooyan AA, Ahmed AA. Reward feedback accelerates motor learning. *J Neurophysiol* 113: 633–646, 2015.
- Nourrit-Lucas D, Zelic G, Deschamps T, Hilpron M, Delignières D. Persistent coordination patterns in a complex task after 10 years delay. *Hum Mov Sci* 32: 1365–78, 2013.
- Park SW, Dijkstra TM, Sternad D. Learning to never forget—time scales and specificity of long-term memory of a motor skill. *Front Comput Neurosci* 7: 111, 2013.
- Park SW, Sternad D. Robust retention of individual sensorimotor skill after self-guided practice. *J Neurophysiol* 113: 2635–2645, 2015.
- Peng CK, Buldyrev SV, Havlin S, Simons M, Stanley HE, Goldberger AL. Mosaic organization of DNA nucleotides. *Phys Rev E* 49: 1685–1689, 1994.
- Picard N, Matsuzaka Y, Strick PL. Extended practice of a motor skill is associated with reduced metabolic activity in M1. *Nat Neurosci* 16: 1340–1347, 2013.
- Reis J, Schambra HM, Cohen LG, Buch ER, Fritsch B, Zarahn E, Celnik PA, Krakauer JW. Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. *Proc Natl Acad Sci USA* 106: 1590–1595, 2009.
- Reisman DS, Wityk R, Silver K, Bastian AJ. Split-belt treadmill adaptation transfers to overground walking in persons poststroke. *Neurorehabil Neural Repair* 23: 735–44, 2009.
- Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP. Neural constraints on learning. *Nature* 512: 423–426, 2014.
- Salmoni A, Schmidt R, Walter C. Knowledge of results and motor learning: a review and critical reappraisal. *Psychol Bull* 95: 355–386, 1984.
- Schmidt RA, Lee TD. *Motor Control and Learning: A Behavioral Emphasis* (5th ed.). Champaign, IL: Human Kinetics, 2011.
- Selen LP, Franklin DW, Wolpert DM. Impedance control reduces instability that arises from motor noise. *J Neurosci* 29: 12606–12616, 2009.
- Shapiro SS, Wilk MB. An analysis of variance test for normality (complete samples). *Biometrika* 52: 591–611, 1965.
- Shmuelof L, Krakauer JW, Mazzoni P. How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J Neurophysiol* 108: 578–594, 2012.
- Smeets JB, Frens MA, Brenner E. Throwing darts: timing is not the limiting factor. *Exp Brain Res* 144: 268–274, 2002.
- Smith MA, Ghazizadeh A, Shadmehr R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4: e179, 2006.
- Sternad D, Abe MO. Variability, noise, and sensitivity to error in learning a motor task. In: *Motor Control: Theories, Experiments, and Applications*, edited by Danion F, Latash M. New York: Oxford Univ. Press, 2010, p. 267–294.

- Sternad D, Abe MO, Hu X, Müller H.** Neuromotor noise, error tolerance and velocity-dependent costs in skilled performance. *PLoS Comput Biol* 7: e1002159, 2011.
- Sternad D, Huber ME, Kuznetsov N.** Acquisition of novel and complex motor skills: stable solutions where intrinsic noise matters less. *Adv Exp Med Biol* 826: 101–124, 2014.
- Theiss RD, Heckman CJ.** Systematic variation in effects of serotonin and norepinephrine on repetitive firing properties of ventral horn neurons. *Neuroscience* 134: 803–815, 2005.
- Thorndike EL.** The law of effect. *Am J Psychol* 39: 212–222, 1927.
- Thoroughman KA, Shadmehr R.** Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19: 8573–8588, 1999.
- Thoroughman KA, Shadmehr R.** Learning of action through adaptive combination of motor primitives. *Nature* 407: 742–747, 2000.
- Tobler PN, O’Doherty JP, Dolan RJ, Schultz W.** Human neural learning depends on reward prediction errors in the blocking paradigm. *J Neurophysiol* 95: 301–310, 2006.
- Todorov E, Jordan MI.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
- Wei K, Glaser JI, Deng L, Thompson CK, Stevenson IH, Wang Q, Hornby TG, Heckman CJ, Kording KP.** Serotonin affects movement gain control in the spinal cord. *J Neurosci* 34: 12690–12700, 2014.
- Wolpert DM, Ghahramani Z, Flanagan JR.** Perspective and problems in motor learning. *Trends Cogn Sci* 5: 487–494, 2001.
- Woodworth RS.** The accuracy of voluntary movement. *Psychol Rev* 3: 1–114, 1899.
- Wu HG, Miyamoto YR, Castro LN, Ölveczky BP, Smith MA.** Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci* 17: 312–321, 2014.

